

WHAT IS CLAIMED IS:

- 1 1. A method of optimizing a query in a multi-tenant database, said
2 database having one or more data tables, each table having one or more logical columns
3 defining data categories and one or more logical rows associated with one or more tenants,
4 wherein a plurality of tenants have data stored in the data tables, the method comprising:
5 generating tenant-level statistics for each of said plurality of tenants for each
6 of the data tables;
7 receiving a SQL query; and
8 optimizing the SQL query based on the tenant-level statistics.
- 1 2. The method of claim 1, wherein each tenant includes one or more
2 associated users, the method further including:
3 generating user-level statistics for each user of each tenant for each of the data
4 tables; and
5 optimizing the SQL query based on the user-level statistics.
- 1 3. The method of claim 2, wherein the user-level statistics are stored to a
2 user metadata table.
- 1 4. The method of claim 2, wherein generating user-level statistics
2 includes determining a total number of distinct rows for each of said plurality of users.
- 1 5. The method of claim 4, wherein the total number is an approximate
2 number based on one or more of a) a number of rows viewable by the user and users below
3 the user in a role hierarchy, b) a number of rows that are shared by a group to which the user
4 belongs and c) a number of rows that are manually shared to the user by another user or
5 group of users.
- 1 6. The method of claim 2, wherein generating user-level statistics for a
2 user is performed according to one of a) on a scheduled basis, b) after a predetermined
3 number of queries by the user, and c) each time an unconstrained query is run by the user.
- 1 7. The method of claim 1, wherein generating tenant-level statistics is
2 performed on a periodic basis.

1 8 The method of claim 1, wherein generating includes determining a
2 total number of distinct rows accessible for each of said plurality of tenants.

1 9. The method of claim 8, wherein the tenant-level statistics are stored to
2 a tenant metadata table.

1 10. The method of claim 1, wherein at least one column of one of said
2 tables includes data associated with two or more tenants.

1 11. A multi-tenant database system, comprising:
2 a database having one or more data tables, each table having one or more
3 columns defining data categories and one or more rows associated with one or more tenants,
4 wherein a plurality of tenants have data stored in the data tables;
5 a statistics generating module configured to generate tenant-level statistics for
6 each tenant for each of the data tables; and
7 a query optimization module, configured to optimize a database query based
8 on the tenant-level statistics.

1 12. The multi-tenant database system of claim 11, wherein each tenant
2 includes one or more associated users, wherein the statistics generating module is further
3 configured to generate user-level statistics for each user, and wherein the query optimization
4 module is further configured to optimize the database query based on the user-level statistics.

1 13. The system of claim 12, further including a memory module, wherein
2 the statistics generating module stores the user-level statistics to a metadata table in the
3 memory module.

1 14. The system of claim 12, wherein the statistics generating module
2 generates user-level statistics by determining a total number of distinct rows for each of said
3 plurality of users.

1 15. The system of claim 14, wherein the total number is an approximate
2 number based on one or more of a) a number of rows viewable by the user and users below
3 the user in a role hierarchy, b) a number of rows that are shared by a group to which the user
4 belongs and c) a number of rows that are manually shared to the user by another user or
5 group of users.

1 16. The system of claim 12, wherein the statistics generating module
2 generates user-level statistics for a user according to one of a) on a scheduled basis, b) after a
3 predetermined number of queries by the user, and c) each time an unconstrained query is run
4 by the user.

1 17. The system of claim 11, wherein the statistics generating module
2 generates tenant-level statistics on a periodic basis.

1 18. The system of claim 11, wherein the statistics generating module
2 generates tenant-level statistics by determining a total number of distinct rows viewable for
3 each of said plurality of tenants.

1 19. The system of claim 18, further including a memory module, wherein
2 the statistics generating module stores the tenant-level statistics to a tenant metadata table in
3 the memory module.

1 20. The system of claim 11, wherein at least one column of one of said
2 tables includes data associated with two or more tenants.

1 21. A method of optimizing a query in a multi-tenant database, said
2 database having one or more data tables, each table having one or more logical columns
3 defining data categories and one or more logical rows associated with one or more tenants,
4 wherein a plurality of tenants have data stored in the data tables, and wherein each tenant
5 includes one or more users, the method comprising:
6 processing the data tables so as to determine tenant-level statistics for each of
7 said plurality of tenants;
8 processing the data tables so as to determine user-level statistics for each of
9 said plurality of user;
10 receiving a SQL query; and
11 optimizing the SQL query based on one or both of the tenant-level statistics
12 and the user-level statistics.

1 22. The method of claim 21, further including:
2 storing the user-level statistics to a user-level metadata table in a memory
3 module; and

4 storing the tenant-level statistics to a tenant-level metadata table in the
5 memory module.

1 23. The method of claim 21, wherein determining user-level statistics
2 includes determining a total number of distinct rows for each of said plurality of users, and
3 wherein determining tenant-level statistics includes determining a total number of distinct
4 rows for each of said plurality of tenants.

1 24. The method of claim 21, wherein processing the data tables to
2 determine tenant-level statistics is performed on a periodic basis.

1 25. The method of claim 21, wherein processing the data tables to
2 determine user-level statistics for a user is performed according to one of a) on a scheduled
3 basis, b) after a predetermined number of queries by the user, and c) each time an
4 unconstrained query is run by the user.